

IN THE CLAIMS:

- 1 1. (Currently Amended) A method for enabling parity declustering in a balanced parity
2 array of a storage system, where an operating system performs the method comprising the
3 steps of:
4 combining a plurality of unbalanced parity stripe arrays, each unbalanced parity
5 stripe array storing an unequal number of parity blocks per disk, to form the balanced
6 parity array, the balanced parity array storing substantially the same number of parity
7 blocks on all disks, each unbalanced parity stripe array and the balanced parity array hav-
8 ing parity blocks on a set of storage devices that are disjoint from a set of storage devices
9 storing data blocks; and
10 distributing assignment of storage devices to parity groups throughout the bal-
11 anced parity array.
- 1 2. (Original) The method of Claim 1 further comprising the step of, after a single or
2 double storage device failure, ensuring that all surviving data storage devices are loaded
3 uniformly during reconstruction of the failed storage device or devices.
- 1 3. (Original) The method of Claim 1 wherein the storage system is a filer.
- 1 4. (Currently Amended) The method of Claim 1 further comprising the steps of:
2 dividing each storage device into blocks; and
3 organizing the blocks into stripes across the devices, wherein each stripe contains
4 data and parity blocks from each of the devices of ~~the~~ a balanced array, where the bal-
5 anced array includes the set of storage devices to store data blocks and the set of storage
6 devices to store parity block.

- 1 5. (Original) The method of Claim 4 wherein the step of distributing comprises the step
2 of selecting patterns of characters representing data storage devices of a stripe to thereby
3 change the association of the data storage devices with parity groups from stripe to stripe
4 of the balanced array.
- 1 6. (Original) The method of Claim 5 wherein the characters are binary numbers.
- 1 7. (Original) The method of Claim 5 wherein the characters are ternary numbers.
- 1 8. (Original) The method of Claim 1 further comprising the steps of:
2 configuring the balanced array as a RAID-4 style array;
3 initially under-populating the array with storage devices; and
4 adding storage devices until a fully populated array of predetermined size is
5 achieved.
- 1 9. (Original) The method of Claim 8 wherein the storage devices are disks.
- 1 10. (Currently Amended) A system that enables parity declustering in a balanced parity
2 array of a storage system, the system comprising:
3 a plurality of storage devices, each storage device divided into blocks that are fur-
4 ther organized into stripes, wherein each stripe contains data and parity blocks from each
5 of the devices of ~~the~~ a balanced array;
6 a storage operating system including a storage layer configured to implement a
7 parity assignment technique that distributes assignment of devices to parity groups
8 throughout the balanced array such that all storage devices contain the same amount of
9 information; and
10 a processing element configured to execute the operating system to thereby in-
11 voke storage access operations to and from the balanced parity array in accordance with a
12 concentrated parity technique, where the concentrated parity technique combines unbal-

13 anced parity stripe arrays from two or more of the parity groups to form a balanced parity
14 array with an equal number of parity blocks per disk storing parity.

1 11. (Currently Amended) The system of Claim 10 wherein ~~the storage layer further~~
2 ~~combines a plurality of unbalanced stripe arrays to form the balanced array,~~ each unbal-
3 anced stripe array having parity blocks on a set of storage devices that are disjoint from a
4 set of storage devices storing data blocks.

1 12. (Original) The system of Claim 11 wherein the storage devices are disks and
2 wherein the storage layer is a RAID layer.

1 13. (Original) The system of Claim 12 wherein the RAID layer is implemented in logic
2 circuitry.

1 14. (Original) The system of Claim 10 wherein the storage system is a network-attached
2 storage appliance.

1 15. (Original) The system of Claim 10 wherein the storage devices are one of video
2 tape, optical, DVD, magnetic tape and bubble memory devices.

1 16. (Original) The system of Claim 10 wherein the storage devices are media adapted to
2 store information contained within the data and parity blocks.

1 17. (Currently Amended) Apparatus for enabling parity declustering in a balanced parity
2 array of a storage system, the apparatus comprising:

3 means for combining a plurality of unbalanced parity stripe arrays, each unbal-
4 anced parity stripe array storing an unequal number of parity blocks per disk, to form the
5 balanced parity array, the balanced parity array storing substantially the same number of
6 parity blocks on all disks, each unbalanced parity stripe array and the balanced parity ar-

7 | ray having parity blocks on a set of storage devices that are disjoint from a set of storage
8 | devices storing data blocks; and
9 | means for distributing assignment of devices to parity groups throughout the bal-
10 | anced parity array.

1 | 18. (Currently Amended) The apparatus of Claim 17 further comprising:
2 | means for dividing each storage device into blocks; and
3 | means for organizing the blocks into stripes across the devices, wherein each
4 | stripe contains data and parity blocks from each of the devices of ~~the~~ a balanced array,
5 | where the balanced array includes the set of storage devices to store data blocks and the
6 | set of storage devices to store parity block.

1 | 19. (Original) The apparatus of Claim 18 wherein the means for distributing comprises
2 | means for selecting patterns of characters representing data storage devices of a stripe to
3 | thereby change the association of the data storage devices with parity groups from stripe
4 | to stripe of the balanced array.

1 | 20. (Currently Amended) A computer readable medium containing executable program
2 | instructions for enabling parity declustering in a balanced parity array of a storage sys-
3 | tem, the executable program instructions comprising program instructions for:
4 | combining a plurality of unbalanced parity stripe arrays, each unbalanced parity
5 | stripe array storing an unequal number of parity blocks per disk, to form the balanced
6 | parity array, the balanced parity array storing substantially the same number of parity
7 | blocks on all disks, each unbalanced parity stripe array and the balanced parity array hav-
8 | ing parity blocks on a set of storage devices that are disjoint from a set of storage devices
9 | storing data blocks; and
10 | distributing assignment of devices to parity groups throughout the balanced parity
11 | array.

1 21. (Currently Amended) The computer readable medium of Claim 20 further compris-
2 ing program instructions for:

3 dividing each storage device into blocks; and
4 organizing the blocks into stripes across the devices, wherein each stripe contains
5 data and parity blocks from each of the devices of ~~the~~ a balanced array, where the bal-
6 anced array includes the set of storage devices to store data blocks and the set of storage
7 devices to store parity block.

1 22. (Original) The computer readable medium of Claim 21 wherein the program instruc-
2 tions for distributing comprises program instructions for selecting patterns of characters
3 representing data storage devices of a stripe to thereby change the association of the data
4 storage devices with parity groups from stripe to stripe of the balanced array.

1 23. -37. (Cancelled)

1 38. (Currently Amended) A method for declustering a parity array having a plurality of
2 storage devices, where an operating system performs the method comprising the steps of:

3 assigning a first plurality of data and parity blocks to a first ~~parity~~ group, where
4 the data blocks are stored on a plurality of storage devices that are disjoint from a plural-
5 ity of storage devices storing parity blocks to form a first parity group, wherein the parity
6 blocks are striped across the plurality of storage devices in the first parity group with an
7 uneven number of parity blocks per each storage device;

8 assigning a second plurality of data and parity blocks to a second ~~parity~~ group,
9 where the data blocks are stored on a plurality of storage devices that are disjoint from a
10 plurality of storage devices storing parity blocks to form a second parity group, wherein
11 the parity blocks are striped across the plurality of storage devices in the second parity
12 group with an uneven number of parity blocks per each storage device, wherein the first
13 and second parity groups being independent from each other and distributed throughout
14 the plurality of storage devices of the parity array; and

15 combining the first parity group and the second parity group to form a balanced
16 parity array, the balanced parity array storing substantially the same number of parity
17 blocks on the storage devices configured to store parity.

1 39. (Cancelled)

1 40. (Currently Amended) A method for declustering a parity array having a plurality of
2 storage devices, where an operating system performs the method comprising the step of:
3 assigning a plurality of ~~data and~~ parity blocks to a plurality of parity groups, the
4 plurality of parity groups being independent from each other and distributed throughout
5 the plurality of storage devices of the parity array; and
6 combining the plurality of parity groups to form a balanced array, the balanced ar-
7 ray storing substantially the same number of blocks on all disks.

1 41. (Currently Amended) A disk array having a declustered parity array, comprising:
2 a plurality of storage devices having a first and second parity group;
3 a first plurality of ~~data and~~ parity blocks assigned to the first parity group;
4 a second plurality of ~~data and~~ parity blocks assigned to the second parity group,
5 the first and second parity groups being independent from each other and distributed
6 throughout the plurality of storage devices of the parity array; and
7 a balanced parity array created by combining the first parity group and the second
8 parity group, the balanced parity array storing substantially the same number of blocks on
9 all disks.

1 42. (Cancelled)

1 43. (Currently Amended) A disk array having a declustered parity array, comprising:
2 a plurality of storage devices having a plurality of parity groups;

3 | a plurality of ~~data and~~ parity blocks assigned to the plurality of parity groups, the
4 | plurality of parity groups being independent from each other and distributed throughout
5 | the plurality of storage devices of the parity array; and

6 | a balanced parity array created by combining the plurality parity groups, the bal-
7 | anced parity array storing substantially the same number of parity blocks on all disks.

1 | 44. (Currently Amended) A disk array having a declustered parity array, comprising:

2 | a plurality of storage devices;

3 | means for assigning a first plurality of ~~data and~~ parity blocks to a first parity
4 | group;

5 | means for assigning a second plurality of ~~data and~~ parity blocks to a second parity
6 | group, the first and second parity groups being independent from each other and distrib-
7 | uted throughout the plurality of storage devices of the parity array; and

8 | means for combining the first parity group and the second parity group to form a
9 | balanced parity array, the balanced parity array storing substantially the same number of
10 | parity blocks on all disks.

1 | 45. (Cancelled)

1 | 46. (Currently Amended) A disk array having a declustered parity array, comprising:

2 | a plurality of storage devices;

3 | means for assigning a plurality of ~~data and~~ parity blocks to a plurality of parity
4 | groups, the plurality of parity groups being independent from each other and distributed
5 | throughout the plurality of storage devices of the parity array; and

6 | means for combining the plurality of parity groups to form a balanced parity ar-
7 | ray, the balanced parity array storing substantially the same number of parity blocks on
8 | all disks.

1 | 47. – 54. (Cancelled)

1 55. (Currently Amended) A computer implemented method for enabling parity declus-
2 tering of a storage system, the method comprising the steps of:
3 providing a first array of storage devices for storing data blocks and parity blocks,
4 the data blocks organized into at least one parity group associated with the parity blocks,
5 the first array storing an unequal number of parity blocks on differing ones of the storage
6 devices;
7 providing a second array of storage devices for storing data blocks and parity
8 blocks, the data blocks organized into at least one parity group associated with the parity
9 blocks, the second array storing an unequal number of parity blocks on differing ones of
10 the storage devices;
11 combining the first and second arrays to form a combined array having substan-
12 tially the same number of parity blocks stored on each storage device of the combined
13 array; and
14 changing the association of data blocks with parity groups in the first array and
15 the second array so that each parity group is associated with data blocks that are distrib-
16 uted substantially uniformly throughout the storage devices that store data blocks in the
17 combined array.

1 56. (Previously Presented) The method of claim 55 further comprising the steps of:
2 organizing the data and parity blocks into stripes across the storage devices.

1 57. (Previously Presented) The method of claim 56 wherein the step of redistributing
2 comprises the step of changing the association of the data storage devices with parity
3 groups from stripe to stripe in the combined array.

1 58. (Previously Presented) The method of claim 57 wherein the step of changing further
2 comprises the step of selecting differing patterns of characters representing data storage
3 devices of a stripe.

1 59. (Previously Presented) The method of claim 58 wherein the characters are binary
2 numbers.

1 60. (Previously Presented) The method of claim 58 wherein the characters are ternary
2 numbers.

1 61. (Previously Presented) The method of claim 55 wherein the storage devices are disk
2 drives.

1 62. (Previously Presented) A computer implemented method for enabling parity declus-
2 tering of a storage array having a plurality of storage devices, the method comprising the
3 steps of:
4 dividing each storage device into blocks;
5 organizing the blocks into a plurality of stripes across the storage devices,
6 wherein each stripe contains data and parity blocks;
7 storing data in data blocks and parity information in parity blocks, the parity
8 blocks storing parity information for a plurality of parity groups; and
9 varying the association of the storage devices to parity groups from stripe to stripe
10 in the storage array such that each parity group is associated with data blocks that are dis-
11 tributed substantially uniformly throughout the storage devices that store data blocks in
12 the storage array.

1 63. (Previously Presented) The method of claim 62 wherein the step of changing com-
2 prises the step of selecting differing patterns of characters representing data storage de-
3 vices of a stripe to thereby change the association of data blocks with parity groups from
4 stripe to stripe of the storage array.

1 64. (Currently Amended) An apparatus for enabling parity declustering of a storage sys-
2 tem, the apparatus comprising:

3 a first array of storage devices for storing data blocks and parity blocks, the data
4 blocks organized into at least one parity group associated with the parity blocks, the first
5 array storing an unequal number of blocks on differing ones of the storage devices;

6 a second array of storage devices for storing data blocks and parity blocks, the
7 data blocks organized into at least one parity group associated with the parity blocks, the
8 second array storing an unequal number of blocks on differing ones of the storage de-
9 vices; and

10 a storage operating system configured to combine the first and second arrays to
11 form a combined array having substantially the same number of blocks stored on each
12 storage device of the combined array, and configured to change the association of data
13 blocks with parity groups in the first array and the second array so that each parity group
14 is associated with data blocks that are distributed substantially uniformly throughout the
15 storage devices that store data blocks in the combined array.

1 65. (Previously Presented) The apparatus of claim 64 wherein each the blocks are organ-
2 ized into stripes across the storage devices.

1 66. (Previously Presented) The apparatus of claim 65 wherein the storage devices are
2 disk drives.

1 67. (Previously Presented) The apparatus of claim 64 wherein the storage devices are
2 one of video tape, optical, DVD, magnetic tape and bubble memory devices.

1 68. (Previously Presented) The apparatus of claim 64 wherein the system is a network-
2 attached storage appliance.

1 69. (Currently Amended) An apparatus for enabling parity declustering of a storage ar-
2 ray having a plurality of storage devices, the system comprising:
3 a storage operating system configured to divide each storage device into blocks
4 and organize the blocks into a plurality of stripes across the storage devices, wherein each
5 stripe contains data and parity blocks and store data in data blocks and parity information
6 in parity blocks, the parity blocks storing parity information for a plurality of parity
7 groups; and
8 the storage operating system further configured to vary the association of the stor-
9 age devices to parity groups from stripe to stripe in the storage array such that, each par-
10 ity group is associated with data blocks that are distributed substantially uniformly
11 throughout the storage devices that store data blocks in the storage array.

1 70. (Previously Presented) The system of claim 68 wherein the storage operating system
2 is configured to select differing patterns of characters representing data storage devices of
3 a stripe to thereby change the association of data blocks with parity groups from stripe to
4 stripe of the storage array.

1 71. (Previously Presented) A system that enables parity declustering of a storage sys-
2 tem, the system comprising:
3 means for providing a first array of storage devices for storing data blocks and
4 parity blocks, the data blocks organized into at least one parity group associated with the
5 parity blocks, the first array storing an unequal number of blocks on differing ones of the
6 storage devices;
7 means for providing a second array of storage devices for storing data blocks and
8 parity blocks, the data blocks organized into at least one parity group associated with the
9 parity blocks, the second array storing an unequal number of blocks on differing ones of
10 the storage devices;

11 means for combining the first and second arrays to form a combined array having
12 substantially the same number of blocks stored on each storage device of the combined
13 array; and

14 means for changing the association of data blocks with parity groups in the first
15 array and the second array so that each parity group is associated with data blocks that are
16 distributed substantially uniformly throughout the storage devices that store data blocks
17 in the combined array.

1 72. (Previously Presented) An apparatus for enabling parity declustering of a storage
2 array having a plurality of storage devices, the system comprising:

3 means for dividing each storage device into blocks;

4 means for organizing the blocks into a plurality of stripes across the storage de-
5 vices, wherein each stripe contains data and parity blocks;

6 means for storing data in data blocks and parity information in parity blocks, the
7 parity blocks storing parity information for a plurality of parity groups; and

8 means for varying the association of the storage devices to parity groups from
9 stripe to stripe in the storage array such that, each parity group is associated with data
10 blocks that are distributed substantially uniformly throughout the storage devices that
11 store data blocks in the storage array.

1 73. (Previously Presented) A computer readable medium containing program instruc-
2 tions for execution on a processor, the executable program instructions comprising pro-
3 gram instructions for:

4 dividing each storage device into blocks;

5 organizing the blocks into a plurality of stripes across the storage devices,
6 wherein each stripe contains data and parity blocks;

7 storing data in data blocks and parity information in parity blocks, the parity
8 blocks storing parity information for a plurality of parity groups; and

9 varying the association of the storage devices to parity groups from stripe to stripe
10 in the storage array such that, each parity group is associated with data blocks that are
11 distributed substantially uniformly throughout the storage devices that store data blocks
12 in the storage array.

1 74. (Previously Presented) The apparatus of claim 17 wherein each unbalanced stripe
2 array has fewer parity blocks per disk than data blocks per disk.

1 75. (Previously Presented) The computer readable medium of claim 20 wherein each un-
2 balanced stripe array has fewer parity blocks per disk than data blocks per disk.

1 76. (Currently Amended) An apparatus for parity declustering in a storage system, the
2 apparatus comprising:

3 a storage operating system configured to combine a plurality of first arrays of
4 storage devices, each first array storing an unequal number of blocks per storage device,
5 to form a second array, the second array storing substantially the same number of parity
6 blocks on ~~all~~ storage devices that store parity blocks; and

7 the storage operating system further configured to redistribute the assignment
8 storage devices to parity groups in the second array so that each storage device will have
9 a substantially equal number of blocks associated with each parity group.

1 77. (Currently Amended) An method for parity declustering in a storage system, the
2 method comprising the steps of:

3 combining a plurality of first arrays of storage devices, each first array storing an
4 unequal number of blocks per storage device, to form a second array, the second array
5 storing substantially the same number of parity blocks on ~~all~~ storage devices that store
6 parity blocks; and

7 | configuring the storage operating system ~~further configured~~ to redistribute the as-
8 signment storage devices to parity groups in the second array so that each storage device
9 will have a substantially equal number of blocks associated with each parity group.

1 78. (Currently Amended) A computer implemented method for enabling parity declus-
2 tering of a storage system, the method comprising the steps of:

3 providing a first array of storage devices for storing data blocks and parity blocks,
4 the data blocks organized into at least one parity group associated with the parity blocks,
5 | the first array storing an unequal number of parity blocks on differing ones of the storage
6 devices;

7 providing a second array of storage devices for storing data blocks and parity
8 blocks, the data blocks organized into at least one parity group associated with the parity
9 | blocks, the second array storing an unequal number of parity blocks on differing ones of
10 the storage devices; and

11 combining the first and second arrays to form a combined array having substan-
12 | tially the same number of parity blocks stored on each storage device of the combined
13 array.

1 79. (Currently Amended) An apparatus for enabling parity declustering of a storage sys-
2 tem, the apparatus comprising:

3 a first array of storage devices for storing data blocks and parity blocks, the data
4 blocks organized into at least one parity group associated with the parity blocks, the first
5 | array storing an unequal number of parity blocks on differing ones of the storage devices;

6 a second array of storage devices for storing data blocks and parity blocks, the
7 data blocks organized into at least one parity group associated with the parity blocks, the
8 | second array storing an unequal number of parity blocks on differing ones of the storage
9 devices; and

10 a storage operating system configured to combine the first and second arrays to
11 | form a combined array having substantially the same number of parity blocks stored on
12 each storage device of the combined array.

1 80. (Currently Amended) A method for enabling a balanced arrangement of a storage sys-
2 tem, the method comprising the steps of:

3 creating a plurality of unbalanced stripe arrangements with each unbalanced stripe
4 | arrangement storing an unequal number of parity blocks per disk in an array of disks;

5 combining the plurality of unbalanced stripe arrangements to form the balanced
6 arrangement, with the balanced arrangement storing substantially the same number of
7 | parity blocks per disk in the array of disks configured to store parity;

8 storing parity blocks across one or more disks in the array of disks to create one or
9 more parity groups;

10 storing data blocks across the remaining disks of the disks in the array with the
11 parity blocks and the data blocks stored on different disks of the array; and

12 assigning storage devices to different parity groups throughout the balanced ar-
13 rangement.

1 81. (Currently Amended) A method for enabling a balanced arrangement of a storage sys-
2 tem, the method comprising the steps of:

3 creating a plurality of unbalanced stripe arrangements with each unbalanced stripe
4 | arrangement storing an unequal number of parity blocks per disk in an array of disks; and

5 combining the plurality of unbalanced stripe arrangements to form the balanced
6 arrangement, with the balanced arrangement storing substantially the same number of
7 | parity blocks per disk in the array of disks configured to store parity.